arXiv:1409.0084v1 [cs.CV] 30 Aug 2014

Kernel Coding: General Formulation and Special Cases

Mehrtash T. Harandi and Mathieu Salzmann

Australian National University, Canberra, ACT 0200, Australia
NICTA*, Locked Bag 8001, Canberra, ACT 2601, Australia

Abstract. Representing images by compact codes has proven beneficial for many visual recognition tasks. Most existing techniques, however, perform this coding step directly in image feature space, where the distributions of the different classes are typically entangled. In contrast, here, we study the problem of performing coding in a high-dimensional Hilbert space, where the classes are expected to be more easily separable. To this end, we introduce a general coding formulation that englobes the most popular techniques, such as bag of words, sparse coding and locality-based coding, and show how this formulation and its special cases can be kernelized. Importantly, we address several aspects of learning in our general formulation, such as kernel learning, dictionary learning and supervised kernel coding. Our experimental evaluation on several visual recognition tasks demonstrates the benefits of performing coding in Hilbert space, and in particular of jointly learning the kernel, the dictionary and the classifier.

Keywords: Kernel methods, sparse coding, visual recognition

1 Introduction

Over the years, **coding** -in its broadest definition- has proven a crucial step in visual recognition systems [1,2]. Many techniques have been investigated, such as bag of words [3,4,5,6,7,8], sparse coding [9,10] and locality-based coding [11,12]. All these techniques follow a similar flow: Given a dictionary of codewords, a query is associated to one or multiple dictionary elements with different weights (*i.e.*, binary or real). These weights, or *codes*, act as the new representation for the query and serve as input to a classifier (*e.g.*, Support Vector Machine (SVM)) after an optional pooling step.

From the description above, it is quite clear that the quality of the codes will depend on the dictionary. Therefore, recent research has focused on learning codebooks that better reflect the underlying recognition problem from training data [13,14,15]. In particular, supervised dictionary learning methods have been introduced to directly exploit the labels of the training samples in the objective function of the coding problem, thus jointly learning a codebook and a classifier [16,17,18,19].

In most existing techniques, coding is performed directly in image feature space. In realistic scenarios, the classes are not linearly separable in this space, and thus kernel-based classifiers are learned on the codes. However, the codes themselves may have

* NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy, as well as by the Australian Research Council through the ICT Centre of Excellence program.

been affected by the intricate distributions of the classes in feature space. Therefore, they may be sub-optimal for classification, even when employing a non-linear classifier. While supervised dictionary learning methods may help computing better-suited codes, they are currently limited to linear, or bilinear classifiers, which may be ill-suited for complex visual recognition tasks.

To alleviate the aforementioned issue, a handful of studies have considered performing coding in kernel space [20,21]. It is widely acknowledged that mapping the data to a high-dimensional Hilbert space should make the classes more easily separable. Therefore, one can reasonably expect that codes extracted in Hilbert space should better reflect the classification problem at hand, while still yielding a compact representation of the data. However, as mentioned earlier, the literature on this topic remains limited, and exclusively focused on the specific case of sparse coding. In [20], Gao *et al*. kernelized the Lasso problem and introduced a gradient descend approach to obtaining a dictionary in kernel space. This approach, however, is restricted to the Gaussian kernel. In [21], Van Nguyen *et al*. addressed the problem of unsupervised dictionary learning in kernel space by kernelizing the KSVD algorithm. While this generalizes kernel dictionary learning beyond the Gaussian kernel case, many questions remain unanswered. For instance:

1. *Is there a principled way to find the best kernel space for the problem at hand?*
2. *How can a supervised dictionary be learned in kernel space?*
3. *Can coding schemes other than sparse coding be performed in the kernel space?*

Here, we aim to answer such questions and therefore study several aspects of coding in a possibly infinite-dimensional Hilbert space. To this end, we first introduce a general formulation to the coding problem, which encompasses bags-of-words, sparse coding and locality-based coding. We show that, in all these cases, coding in a Hilbert space can be fully expressed in terms of kernels. We then show how the kernel, the dictionary, and the classifier can be learned in our general framework. This allows us to find the Hilbert space, codes and classifier that are jointly best-suited for the task at hand.

We evaluate our general kernel coding formulation on several image recognition benchmarks and provide a detailed analysis of its different special cases. Our empirical results evidence the strengths of our kernel coding approach over linear methods. Furthermore, it shows the benefits of jointly learning the kernel, dictionary and classifier over existing kernel sparse coding methods that only learn the codebook.

2 Kernel Coding

Given a query vector $\boldsymbol{x} \in \mathbb{R}^d$, such as an image descriptor, many methods have been proposed to transform $\boldsymbol{x}$ into a more compact, and hopefully discriminative, representation, hereafter referred to as *code*. Popular examples of such an approach include sparse coding and bag of words. In the following, we introduce a general formulation that englobes many different coding techniques, and lets us derive kernelized versions of these different methods.

More specifically, given a dictionary $\boldsymbol{D}_{d\times N} = [\boldsymbol{d}_1|\boldsymbol{d}_2|\cdots|\boldsymbol{d}_N]$, $\boldsymbol{d}_i \in \mathbb{R}^d$ containing N elements, coding can be expressed as the solution to the optimization problem

$$\begin{aligned} &\min_{\boldsymbol{y}} \left\|\boldsymbol{x} - \sum_{j=1}^{N} [\boldsymbol{y}]_j \boldsymbol{d}_j\right\|_2^2 + \gamma r(\boldsymbol{y}; \boldsymbol{x}, \boldsymbol{D}) \\ &\text{s.t.} \quad \boldsymbol{y} \in C, \end{aligned} \tag{1}$$

where $[\cdot]_j$ denotes the j^{th} element of a vector, $r(\cdot)$ is a prior on the codes $\boldsymbol{y}$, and C is a set of constraints on $\boldsymbol{y}$. Note that this formulation allows the prior to be dependent on both the query $\boldsymbol{x}$ and the dictionary $\boldsymbol{D}$. Although not explicit, this is also true for C.

Our goal here is to perform nonlinear coding in the hope to obtain a better representation. To this end, let $\phi : \mathbb{R}^d \to \mathcal{H}$ be a mapping to a Reproducing Kernel Hilbert Space (RKHS) induced by the kernel $k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \phi(\boldsymbol{x}_i)^T \phi(\boldsymbol{x}_j)$. Then, coding in $\mathcal{H}$ can be formulated as

$$\begin{aligned} &\min_{\boldsymbol{y}} \left\|\phi(\boldsymbol{x}) - \sum_{j=1}^{N} [\boldsymbol{y}]_j \phi(\boldsymbol{d}_j)\right\|_2^2 + \gamma r(\boldsymbol{y}; \phi(\boldsymbol{x}), \phi(\boldsymbol{D})) \\ &\text{s.t.} \quad \boldsymbol{y} \in C. \end{aligned} \tag{2}$$

Expanding the first term in (2) yields

$$\begin{aligned} \left\|\phi(\boldsymbol{x}) - \sum_{j=1}^{N} [\boldsymbol{y}]_j \phi(\boldsymbol{d}_j)\right\|_2^2 &= \phi(\boldsymbol{x})^T \phi(\boldsymbol{x}) - 2\sum_{j=1}^{N} [\boldsymbol{y}]_j \phi(\boldsymbol{d}_j)^T \phi(\boldsymbol{x}) + \sum_{i,j=1}^{N} [\boldsymbol{y}]_i [\boldsymbol{y}]_j \phi(\boldsymbol{d}_i)^T \phi(\boldsymbol{d}_j) \\ &= k(\boldsymbol{x}, \boldsymbol{x}) - 2\sum_{j=1}^{N} [\boldsymbol{y}]_j k(\boldsymbol{x}, \boldsymbol{d}_j) + \sum_{i,j=1}^{N} [\boldsymbol{y}]_i [\boldsymbol{y}]_j k(\boldsymbol{d}_i, \boldsymbol{d}_j) \\ &= k(\boldsymbol{x}, \boldsymbol{x}) - 2\boldsymbol{y}^T \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{D}) + \boldsymbol{y}^T \boldsymbol{K}(\boldsymbol{D}, \boldsymbol{D}) \boldsymbol{y}, \end{aligned} \tag{3}$$

where $\boldsymbol{k}(\boldsymbol{x}, \boldsymbol{D}) \in \mathbb{R}^{N\times 1}$, and $\boldsymbol{K}(\boldsymbol{D}, \boldsymbol{D}) \in \mathbb{R}^{N\times N}$. This shows that the reconstruction term in (1), common to most coding techniques, can be kernelized. More importantly, after kernelization, this term remains quadratic, convex and similar to its counterpart in Euclidean space. In the remainder of this section, we discuss the priors $r(\cdot)$ and constraint sets C corresponding to specific coding techniques, and show how they can also be kernelized.

2.1 Kernel Bag of Words

Bag of words [22] is possibly the simplest coding technique, which consists of assigning the query $\boldsymbol{x}$ to a single dictionary element. In our framework, this can be expressed by having no prior $r(\cdot)$ and defining the constraint set C as

$$C = \left\{ \boldsymbol{y} \mid \boldsymbol{y} \in \{0, 1\}^N, \sum_i \boldsymbol{y}_i = 1 \right\}. \tag{4}$$

Note that this constraint set only depends on the code $\boldsymbol{y}$. Therefore, kernelizing this model does not require any other operation than that described in Eq. 3. In other words, $\boldsymbol{y}$ can be obtained by finding the nearest neighbor to $\boldsymbol{x}$ in kernel space and setting the corresponding entry in $\boldsymbol{y}$ to 1.

The above-mentioned hard assignment bag of words model was relaxed to soft assignments [7], where the query is assigned to multiple dictionary elements, with weights depending on the distance between these elements and the query. In this soft assignment bag of words model, a code $\boldsymbol{y}_i$ is directly constrained to take the value

$$\boldsymbol{y}_i = \frac{\exp\left(-\sigma\|\boldsymbol{x} - \boldsymbol{d}_i\|_2^2\right)}{\sum_j \exp\left(-\sigma\|\boldsymbol{x} - \boldsymbol{d}_j\|_2^2\right)} \ . \tag{5}$$

Here, $\sigma > 0$ is the bandwidth parameter of the algorithm. To kernelize this model, we note that each exponential term can be written as

$$\exp\left(-\sigma\|\phi(\boldsymbol{x}) - \phi(\boldsymbol{d}_i)\|_2^2\right) = \exp\left(-\sigma\left(k(\boldsymbol{x},\boldsymbol{x}) - 2k(\boldsymbol{x},\boldsymbol{d}_i) + k(\boldsymbol{d}_i,\boldsymbol{d}_i)\right)\right) \ . \tag{6}$$

Since this only depends on kernel values, so do the resulting codes $\boldsymbol{y}_i$. In practice, soft assignment bags of words have often proven more effective than hard assignment ones, especially when a single code is used to represent an entire image.

2.2 Kernel Sparse Coding

Over the years, sparse coding has proven effective at producing compact and discriminative representations of images [23]. From (1), sparse coding can be obtained by employing the prior

$$r(\boldsymbol{y}) = \|\boldsymbol{y}\|_1 \ , \tag{7}$$

which corresponds to a convex relaxation of the ℓ_0-norm encoding sparsity, and not using any constraints. Since this prior only depends on $\boldsymbol{y}$, the only step required to kernelize sparse coding is given in Eq. 3. Note that any structured, or group, sparsity prior can also be utilized in the same manner.

A solution to kernel sparse coding can be obtained by transforming the resulting optimization problem into a standard vectorized sparse coding problem. To this end, let $\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^T$ be the SVD of $\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D})$. Then it can easily be verified that kernel sparse coding is equivalent to the optimization problem

$$\min_{\boldsymbol{y}} \|\tilde{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{y}\|_2^2 + \gamma\|\boldsymbol{y}\|_1, \tag{8}$$

where $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}^{1/2}$ and $\tilde{\boldsymbol{x}} = \boldsymbol{U}\boldsymbol{\Sigma}^{-1/2}\boldsymbol{k}(\boldsymbol{x},\boldsymbol{D})$. As a consequence, any efficient sparse solver such as SLEP [24] or SPAMS [25] can be employed to solve kernel sparse coding.

2.3 Locality-Constrained Kernel Coding

As with bags of words, the notion of locality between the query and the dictionary elements has also been employed in the sparse coding framework. This was first introduced in the Local Coordinate Coding (LCC) model [11], which was then modified as

the Locality-Constrained Linear Coding (LLC) model [12], whose solution can be obtained more efficiently. We therefore focus on kernelizing the LLC formulation, though it can easily be verified that LCC can be kernelized in a similar manner. From our general formulation (1), LLC can be obtained by defining

$$r(\boldsymbol{y};\boldsymbol{x},\boldsymbol{D}) = \|\boldsymbol{E}\boldsymbol{y}\|_2^2\,,\ \text{with}\ \boldsymbol{E}_{ii} = \exp\left(\sigma\|\boldsymbol{x}-\boldsymbol{d}_i\|_2\right),\ \text{and}\ \ \boldsymbol{E}_{ij} = 0,\ i \neq j\,, \tag{9}$$

and $C = \{\boldsymbol{y} \mid \sum_i \boldsymbol{y}_i = 1\}$. Since the constraints only depend on the codes, they do not require any special care for kernelization. To kernelize the prior, we note that

$$\boldsymbol{E}_{ii} = \exp\left(-\sigma\,\|\phi(\boldsymbol{x}) - \phi(\boldsymbol{d}_i)\|_2\right) = \exp\left(-\sigma\,\sqrt{k(\boldsymbol{x},\boldsymbol{x}) - 2k(\boldsymbol{x},\boldsymbol{d}_i) + k(\boldsymbol{d}_i,\boldsymbol{d}_i)}\right). \tag{10}$$

Therefore, $\boldsymbol{E}$ can be expressed solely in terms of kernel values, and so does $r(\boldsymbol{y};\boldsymbol{x},\boldsymbol{D})$.

In [12], the initial LLC formulation was then approximated to improve the speed of coding. To this end, for each query, the dictionary $\boldsymbol{D}$ was replaced by a local dictionary $\boldsymbol{B}$ formed by the $N_{\boldsymbol{B}}$ nearest dictionary elements to $\boldsymbol{x}$. In Hilbert space, we can follow a similar idea and obtain a local dictionary $\boldsymbol{B}$ by performing kernel nearest neighbor between the dictionary elements and the query. This lets us write LLC in Hilbert space as

$$\begin{aligned} &\min_{\boldsymbol{y}} \|\phi(\boldsymbol{x}) - \phi(\boldsymbol{B})\boldsymbol{y}\|_2^2 \\ &\text{s.t.}\quad \boldsymbol{1}^T\boldsymbol{y} = 1, \end{aligned} \tag{11}$$

which has a form similar to (2) with no prior. This can then be directly kernelized by making use of Eq. 3.

The solution to kernel LLC can be obtained from the Lagrange dual [26] of Eq. (11), which can be written as

$$L_{kllc}(\boldsymbol{y},\lambda) = k(\boldsymbol{x},\boldsymbol{x}) - 2\boldsymbol{y}^T\boldsymbol{k}(\boldsymbol{x},\boldsymbol{D}) + \boldsymbol{y}^T\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D})\boldsymbol{y} + \lambda(\boldsymbol{y}^T\boldsymbol{1} - 1)\,. \tag{12}$$

The gradient of $L_{kllc}(\boldsymbol{y},\lambda)$ w.r.t. $\boldsymbol{y}$ can then be expressed as

$$\begin{aligned} \nabla_{\boldsymbol{y}} L_{kllc} &= -2\boldsymbol{k}(\boldsymbol{x},\boldsymbol{D}) + 2\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D})\boldsymbol{y} + \lambda\boldsymbol{1} \\ &= -2\sum\nolimits_j \boldsymbol{y}_j\boldsymbol{k}(\boldsymbol{x},\boldsymbol{D}) + 2\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D})\boldsymbol{y} + \lambda\boldsymbol{1} \\ &= -2\left(\boldsymbol{1}^T \otimes \boldsymbol{k}(\boldsymbol{x},\boldsymbol{D})\right)\boldsymbol{y} + 2\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D})\boldsymbol{y} + \lambda\boldsymbol{1}\,, \end{aligned} \tag{13}$$

where the second line was obtained by making use of the constraint on $\boldsymbol{y}$. By setting this gradient to 0, $\boldsymbol{y}$ can be obtained as the solution to the linear system $(\boldsymbol{K}(\boldsymbol{D},\boldsymbol{D}) - (\boldsymbol{1}^T \otimes \boldsymbol{k}(\boldsymbol{x},\boldsymbol{D})))\boldsymbol{y} = \boldsymbol{1}$, and then normalized by its ℓ_1 norm to satisfy the constraint[1].

3 Learning in Kernel Coding

In this section, we discuss the learning of the different components of the kernel coding methods described in Section 2. In particular, given a set of M training samples

[1] Note that the dependency on λ is ignored in the linear system since it would only influence the global scale of $\boldsymbol{y}$, which would then be canceled out by the normalization.

$\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^M$, we investigate how the kernel, as well as the dictionary can be learned. Furthermore, we show how a discriminative classifier can be directly introduced and learned within our kernel coding framework. While, in the following, we consider the different learning problems separately, they can of course be solved jointly by employing the alternating minimization approach commonly used in dictionary learning.

3.1 Kernel Learning

We first tackle the problem of finding the best kernel for the problem at hand. While the kernel could, in principle, be considered as a completely free entity, this typically yields an under-constrained formulation and only suits the transductive settings. Here instead, we assume that the kernel has either a parametric form (*e.g.*, Gaussian kernel with width as a parameter), or can be represented as a linear combination of fixed kernels (*i.e.*, multiple kernel learning).

Learning Kernel Parameters

Let us assume that we are given a fixed dictionary $\boldsymbol{D}$, and that our kernel function has a parametric form with parameters β. Following our general formulation, β can be learned by solving the optimization problem

$$\begin{aligned} &\min_{\beta,\{\boldsymbol{y}_i\}} \frac{1}{M}\sum_{i=1}^{M} L_\phi(\beta, \boldsymbol{y}_i; \boldsymbol{x}_i, \boldsymbol{D}) \\ &\text{s.t.} \quad \boldsymbol{y}_i \in C, \ \forall i \in [1, M], \end{aligned} \tag{14}$$

where $\boldsymbol{y}_i$ is the vector of sparse codes for the i^{th} training sample $\boldsymbol{x}_i$, and $L_\phi(\cdot)$ is the kernelized objective function defined in (2).

Note that (14) is not jointly convex in β and $\{\boldsymbol{y}_i\}_{i=1}^M$. Therefore, we follow the standard alternating minimization strategy that consists of iteratively fixing one variable (*i.e.*, either β, or the $\boldsymbol{y}_i$s) and solving for the other. In general, we cannot expect the objective function to be convex in β, even with fixed $\{\boldsymbol{y}_i\}$. Therefore, we utilize a gradient-based trust-region method to obtain β at each iteration.

While any kernel function could be employed in this framework, in practice, we make use of the Gaussian kernel $k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \exp(-\beta\|\boldsymbol{x}_i, \boldsymbol{x}_j\|_2^2)$. Unfortunately, with this kernel, (14) is ill-posed in terms of β. More specifically, $\beta = 0$ is a minimum of (14). This is due to the fact that, if $\beta \to 0$, all samples in the induced Hilbert space $\mathcal{H}$ collapse to one point, *i.e.*, $\|\phi(\boldsymbol{x}_i) - \phi(\boldsymbol{x}_j)\|^2 = k(\boldsymbol{x}_i, \boldsymbol{x}_i) - 2k(\boldsymbol{x}_i, \boldsymbol{x}_j) - k(\boldsymbol{x}^j, \boldsymbol{x}^j) = 0$[2].

To avoid this trivial solution, we propose to search for a β that not only minimizes the kernel coding cost, but also maximizes a measure of discrepancy between the dictionary atoms in $\mathcal{H}$. In other words, we search for a Hilbert space $\mathcal{H}$ that simultaneously yields a diverse dictionary and a good representation of the data. To this end, we define the discrepancy between the dictionary atoms in $\mathcal{H}$ as

$$J_\phi(\boldsymbol{D}, \beta) = \frac{1}{N^2}\sum_{(i,j)=1}^{N} \|\phi(\boldsymbol{d}_i) - \phi(\boldsymbol{d}_j)\|_2^2 = \frac{1}{N^2}\sum_{(i,j)=1}^{N} \Big(k(\boldsymbol{d}_i, \boldsymbol{d}_i) - 2k(\boldsymbol{d}_i, \boldsymbol{d}_j) + k(\boldsymbol{d}_j, \boldsymbol{d}_j)\Big). \tag{15}$$

[2] The same statement holds for polynomial kernels of the form $k(\boldsymbol{x}_i, \boldsymbol{x}_j) = (1 + \beta \boldsymbol{x}_i^T \boldsymbol{x}_j)^p$.

Given $\{\boldsymbol{y}_i\}$, this lets us obtain β by solving the optimization problem

$$\min_{\beta} \frac{\frac{1}{M}\sum_{i=1}^{M} L_\phi(\beta, \boldsymbol{y}_i; \boldsymbol{x}_i, \boldsymbol{D})}{J_\phi(\boldsymbol{D}, \beta)} \quad \text{s.t.} \quad \boldsymbol{y}_i \in C, \ \forall i \in [1, M] \tag{16}$$

We obtain a local minimum to this problem using a gradient-based trust-region method. Note that, in most of the special cases discussed in Section 2, the prior $r(\cdot)$ does not depend on the kernel and can thus be omitted when updating β. This is also the case of the constraint set C, except for soft assignment kernel bag of words where the constraints directly define the codes. In this case, however, the codes can be replaced by the form in Eq. 6 in the objective function, which then becomes a function of β solely.

Multiple Kernel Coding

In various applications, combining multiple kernels has often proven more effective than using a single kernel [27]. Following this idea, we propose to model our kernel $k(\cdot, \cdot)$ as a linear combination of a set of kernels, *i.e.*, $k(\boldsymbol{x}_i, \boldsymbol{x}_j) = \sum_{l=1}^{L} \alpha_l k_l(\boldsymbol{x}_i, \boldsymbol{x}_j)$. Kernel learning then boils down to finding the best coefficients $\boldsymbol{\alpha}$. In our context, this can be done by solving the optimization problem

$$\min_{\boldsymbol{\alpha}, \{\boldsymbol{y}_i\}} \frac{\frac{1}{M}\sum_{i=1}^{M} L_\phi(\boldsymbol{\alpha}, \boldsymbol{y}_i; \boldsymbol{x}_i, \boldsymbol{D})}{J_\phi(\boldsymbol{D}, \boldsymbol{\alpha})} \quad \text{s.t.} \quad \boldsymbol{y}_i \in C, \ \forall i \in [1, M], \tag{17}$$

where $L_\phi(\cdot)$ and $J_\phi(\cdot)$ are defined as before, but in terms of a kernel given as the linear combination of multiple base kernels with weights $\boldsymbol{\alpha}$. As in the parametric case, we obtain $\boldsymbol{\alpha}$ using a gradient-based trust-region method.

3.2 Dictionary Learning

In many cases, it is beneficial to not only learn codes for a given dictionary, but optimize the dictionary to best suit the problem at hand. Here, we show how this can be done in our general formulation. Given fixed kernel parameters and codes for the training data, learning the dictionary can be expressed as solving the optimization problem

$$\min_{\boldsymbol{D}} \frac{1}{M}\sum_{i=1}^{M} L_\phi(\boldsymbol{D}; \boldsymbol{y}_i, \boldsymbol{x}_i) \quad \text{s.t.} \quad \boldsymbol{y}_i \in C, \ \forall i \in [1, M]. \tag{18}$$

Here, we make use of the *Representer theorem* [28] which enables us to express the dictionary as a linear combination of the training samples in RKHS. That is

$$\phi(\boldsymbol{d}_r) = \sum_{i=1}^{M} a_{r,i} \phi(\boldsymbol{x}_i), \tag{19}$$

where $\{a_{r,i}\}$ is the set of weights, now corresponding to our new unknowns. By stacking these weights for the N dictionary elements and the M training samples in a matrix $\boldsymbol{A}_{M\times N} = [\boldsymbol{a}_1|\boldsymbol{a}_2|\cdots|\boldsymbol{a}_N]$, with $\boldsymbol{a}_r = [a_{r,1}, a_{r,2}, \cdots, a_{r,M}]^T$, we have

$$\boldsymbol{\Phi}(\boldsymbol{D}) = \boldsymbol{\Phi}(\boldsymbol{X})\boldsymbol{A} \ . \tag{20}$$

For most of the algorithms presented in Section 2, *i.e.*, kernel bag of words, kernel sparse coding and approximate kernel LLC, the only term that depends on the dictionary is the reconstruction error (the first term in the objective of (2)). Given the matrix of sparse codes $\boldsymbol{Y}_{N\times M}$, this term can be expressed as

$$\begin{aligned} R(\boldsymbol{A}) &= \left\|\boldsymbol{\Phi}(\boldsymbol{X}) - \boldsymbol{\Phi}(\boldsymbol{X})\boldsymbol{A}\boldsymbol{Y}\right\|_F^2 \\ &= \mathrm{Tr}\left(\boldsymbol{\Phi}(\boldsymbol{X})(\mathbf{I}_M - \boldsymbol{A}\boldsymbol{Y})(\mathbf{I}_m - \boldsymbol{A}\boldsymbol{Y})^T\boldsymbol{\Phi}(\boldsymbol{X})^T\right) \\ &= \mathrm{Tr}\left(\boldsymbol{K}(\boldsymbol{X},\boldsymbol{X})(\mathbf{I}_M - \boldsymbol{A}\boldsymbol{Y} - \boldsymbol{Y}^T\boldsymbol{A}^T + \boldsymbol{A}\boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{A}^T)\right) \ . \end{aligned} \tag{21}$$

The new dictionary can then be obtained by zeroing out the gradient of this term w.r.t. $\boldsymbol{A}$. This yields

$$\begin{aligned} \nabla R(\boldsymbol{A}) = 0 &\Rightarrow 2\boldsymbol{Y} = 2\boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{A} \\ &\Rightarrow \boldsymbol{A} = (\boldsymbol{Y}\boldsymbol{Y}^T)^{-1}\boldsymbol{Y} = \boldsymbol{Y}^\dagger \ . \end{aligned} \tag{22}$$

In the case of approximate kernel LLC, we update the full dictionary at once. To this end, for each training sample i, we simply set to 0 the codes corresponding to the elements that do not belong to the local dictionary $\boldsymbol{B}_i$. For soft-assignment kernel bag of words, where the constraints depend on the dictionary, the update of Eq. 22 is not valid and one must resort to an iterative gradient-based update of $\boldsymbol{A}$.

3.3 Supervised Kernel Coding

In the context of sparse coding, several works have studied the idea of making the sparse codes discriminative, and thus jointly learn a classifier with the codes and the dictionary [16,19]. Here, we introduce two formulations that also make use of supervised data in our general kernel coding framework. To this end, given data belonging to S different classes, let $\boldsymbol{l}_i$ be the S-dimensional binary vector encoding the label of training sample $\boldsymbol{x}_i$, *i.e.*, $[\boldsymbol{l}_i]_j = 1$ is sample i belongs to class j.

In our first formulation, we make use of a linear classifier acting on the codes. The prediction of such a classifier takes the form $\hat{\boldsymbol{l}} = \boldsymbol{W}\boldsymbol{y}$. By employing the square loss, learning in Hilbert space can then be written as

$$\begin{aligned} &\min_{\boldsymbol{W},\boldsymbol{D},\{\boldsymbol{y}_i\}} \frac{1}{M}\sum_{i=1}^{M} L_\phi(\boldsymbol{D},\boldsymbol{y}_i;\boldsymbol{x}_i) + \frac{\eta}{M}\sum_{i=1}^{M}\left\|\boldsymbol{l}_i - \boldsymbol{W}\boldsymbol{y}_i\right\|_2^2 + \rho\|\boldsymbol{W}\|_F^2 \ , \\ &\text{s.t.} \quad \boldsymbol{y}_i \in C, \ \forall i \in [1, M], \end{aligned} \tag{23}$$

where we utilize a simple regularizer on the parameters $\boldsymbol{W}$. Following an alternating minimization approach, given $\boldsymbol{D}$ and $\{\boldsymbol{y}_i\}$, the classifier parameters $\boldsymbol{W}$ can be obtained by solving a linear system arising from zeroing out the gradient of the second and third terms in the objective function. While the update for the dictionary is unaffected by the discriminative term, computing the codes must be modified accordingly. However, the least-squares form of this term makes it easy to introduce into the solutions of kernel sparse coding and approximate kernel LLC.

For our second formulation, we employ a bilinear classifier, as suggested in [16]. For a single class j, the prediction of this classifier is given by $\hat{\boldsymbol{l}}_j = \phi(\boldsymbol{x}^i)^T \boldsymbol{W}_j \boldsymbol{y}^i$, which requires one parameter matrix per class. Using the same loss as before, learning can then be expressed as

$$\min_{\{\boldsymbol{W}_j\},\boldsymbol{D},\{\boldsymbol{y}_i\}} \frac{1}{M}\sum_{i=1}^{M} L_\phi(\boldsymbol{D},\boldsymbol{y}_i;\boldsymbol{x}_i) + \frac{\eta}{M}\sum_{i=1}^{M}\sum_{j=1}^{S}\left([\boldsymbol{l}_i]_j - \phi(\boldsymbol{x}_i)^T\boldsymbol{W}_j\boldsymbol{y}_i\right)^2 + \frac{\rho}{S}\sum_{j=1}^{S}\|\boldsymbol{W}_j\|_F^2 \qquad (24)$$

$$\text{s.t.} \quad \boldsymbol{y}_i \in \mathcal{S},\ \forall i \in [1, M].$$

By making use of the Representer theorem, we can express each parameter matrix $\boldsymbol{W}_j$ as a linear combination of the training samples in Hilbert space, which yields

$$\boldsymbol{W}_j = \Phi(\boldsymbol{X})\boldsymbol{A}_j\ . \qquad (25)$$

With this form, the prediction of the classifier is now given by

$$\hat{\boldsymbol{l}}_j = \phi(\boldsymbol{x}_i)^T\Phi(\boldsymbol{X})\boldsymbol{A}_j\boldsymbol{y}_i = \boldsymbol{k}(\boldsymbol{x}_i,\boldsymbol{X})\boldsymbol{A}_j\boldsymbol{y}_i\ , \qquad (26)$$

which depends on the kernel function. Similarly, it can easily be checked that the regularizer on the parameters $\boldsymbol{W}_j$ can be expressed in terms of kernel values. Learning can then be formulated as a function of the matrices $\{\boldsymbol{A}_j\}$, which appear in two convex quadratic terms, and can thus be computed as the solution to a linear system obtained by zeroing out the gradient of the objective function. Similarly to our first formulation, the dictionary update is unaffected by the discriminative terms, and the computation of the codes must account for the additional quadratic terms.

For both formulations, at test time, given a new sample $\boldsymbol{x}^*$, we first compute the codes $\boldsymbol{y}^*$ using the chosen coding approach from Section 2, and then predict the label of the sample by applying the learned classifier.

4 Experimental Evaluation

In this section, we demonstrate the strength of kernel coding on the tasks of material categorization, scene classification, and object and face recognition. We also demonstrate the ability of kernel coding to performing recognition on a Riemannian manifold whose non-linear geometry makes linear coding techniques inapplicable.

Throughout our experiments, we refer to the different coding techniques as

- **kBOW**: kernel (soft-assignment) bag of words,
- **kSC**: kernel sparse coding (8),

- **kLLC**: (approximate) kernel LLC (11),
- **l-SkSC**: supervised kernel sparse coding with the linear classifier from (23),
- **l-SkLLC**: supervised (approximate) kernel LLC with the linear classifier from (23),
- **bi-SkSC**: supervised kernel sparse coding with the bilinear classifier from (24), and
- **bi-SkLLC**: supervised (approximate) kernel LLC with the bilinear classifier from (24).

When not learning a dictionary, classification is performed as proposed in [32]. More specifically, given a query $\boldsymbol{x}$ and its code $\boldsymbol{y}$, let $\boldsymbol{y}_{(s)} = ([\boldsymbol{y}]_1\delta(\gamma_1 - s),\ [\boldsymbol{y}]_2\delta(\gamma_2 - s),\ \cdots,\ [\boldsymbol{y}]_N\delta(\gamma_N - s))^T$ be the class-specific sparse codes for class s, where $\gamma_j \in \mathcal{S}$ is the class label of atom $\boldsymbol{d}_j$ and $\delta(\cdot)$ is the discrete Dirac function. The residual error of query $\boldsymbol{x}$ for class s can be defined as

$$\varepsilon_s(\boldsymbol{x}) = \left\| \boldsymbol{\phi}(\boldsymbol{x}) - \sum_{j=1}^{N} [\boldsymbol{y}]_j \boldsymbol{\phi}(\boldsymbol{d}_j)\delta(\gamma_j - s) \right\|^2 = -2\boldsymbol{y}_{(s)}^T \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{D}) + \boldsymbol{y}_{(s)}^T \boldsymbol{K}(\boldsymbol{D}, \boldsymbol{D})\boldsymbol{y}_{(s)}\ , \qquad (27)$$

where the term $k(\boldsymbol{x}, \boldsymbol{x})$ was dropped since it does not depend on the class label. The label of $\boldsymbol{x}$ is then chosen as the class with minimum residual error $\varepsilon_s(\boldsymbol{x})$. In the case of unsupervised dictionary learning, since the dictionary elements do not have any associated labels anymore, we make use of a simple nearest-neighbor classifier based on the learned codes. Finally, in the supervised scenario, we first obtain the codes and then employ the learned classifier to obtain the labels.

In the following, we used the Gaussian kernel as base kernel, unless stated otherwise. The width of the Gaussian was learned using the method described in Section 3.1 except for the multiple kernel learning experiments. We emphasize that, throughout this section, coding is performed at image level, as in, *e.g.*, [19]. In other words, we represent each image with one descriptor to which a coding technique is applied.

4.1 Flicker Material

As a first experiment, we tackled the problem of material recognition using the Flicker Material Database (FMD) [29], which contains ten categories of materials (*i.e.*, *fabric*, *foliage*, *glass*, *leather*,*metal*, *paper*, *plastic*, *stone*, *water* and *wood*). Each category comprises 100 images, split into 50 close-up views and 50 object-scale views (see Figure. 1). To describe each image, we used a bag of words model with 800 atoms computed from RootSIFT features [31] sampled every 4 pixels. Following the protocol used in [30], half of the images from each category was randomly chosen for training and the rest was used for testing. We report the average classification accuracy over 10 such random partitions.

We first compare the performance of kBOW, kSC, kLLC and linear coding techniques (SC and LLC) without dictionary learning. To this end, we employed all the training samples as dictionary atoms. In Table 1, we report the performance of the different methods on FMD. Note that, with the exception of kBOW, kernel coding techniques significantly outperform the linear ones. In particular, the gap between LLC and its kernel version reaches 4%.

Using the same data, and still without learning a dictionary, we evaluated the performance of multiple kernel coding as described in Section 3.1. To this end, we combined

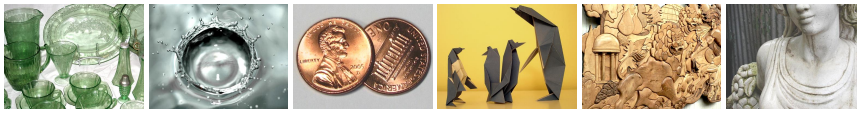

Fig. 1: Examples from the FMD dataset [29]

Method	SC	LLC	kBOW	kSC	kLLC
Accuracy	51.6% ± 2.2	48.8% ± 1.3	46.1% ± 2.0	**53.4**% ± **1.7**	52.8% ± 1.6

Table 1: Comparison of different coding techniques on FMD without dictionary learning.

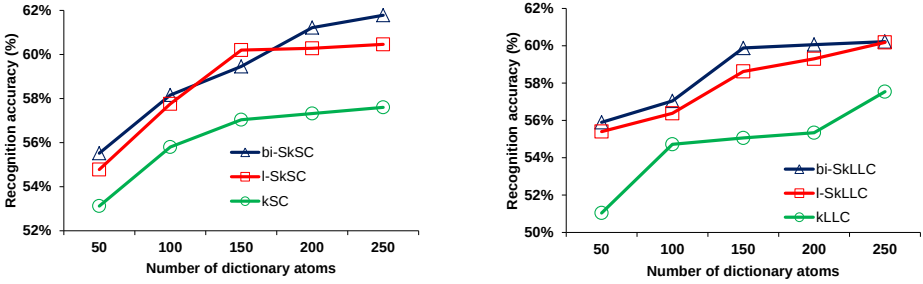

Fig. 2: Recognition accuracies of dictionary learning for unsupervised and supervised kernel coding on FMD.

4 different kernels: a linear kernel, a Gaussian kernel, a second order polynomial kernel and a Sigmoid kernel. After learning the combination, kSC and kLLC achieved **59.1**% and **58.5**%, respectively. This shows that combining kernels can even further improve the results of kernel coding techniques. Note that, in this manner, multiple kernel coding also allows us to nicely combine different features.

We then turn to the cases of dictionary learning for unsupervised and supervised kernel coding. In the unsupervised cases, the dictionary was obtained using the method described in Section 3.2. For the supervised methods, we made us of the technique described in Section 3.3. Fig. 2 depicts the accuracy of the different coding methods as a function of the number of dictionary elements on the FMD dataset. Note that supervised coding consistently improves the classification accuracy for both sparse coding and LLC. Note also that the bilinear classifier outperforms the linear one in most cases. The best recognition accuracy of 61.8% is achieved by supervised kernel sparse coding using a bilinear classifier. To the best of our knowledge, this constitutes the state-of-the-art performance on FMD (Sharan *et al*. [29] reported 60.6% with feature combination). Note that the best performance reported in [29] using only SIFT features is 41.2%, which is significantly lower than kernel coding with ultimately the same features.

4.2 MIT Indoor

As a second experiment, we performed scene classification using the challenging MIT Indoor dataset [33]. MIT Indoor contains 67 scene categories with images collected from online sources such as Flicker and Google. We used a bag of words model with

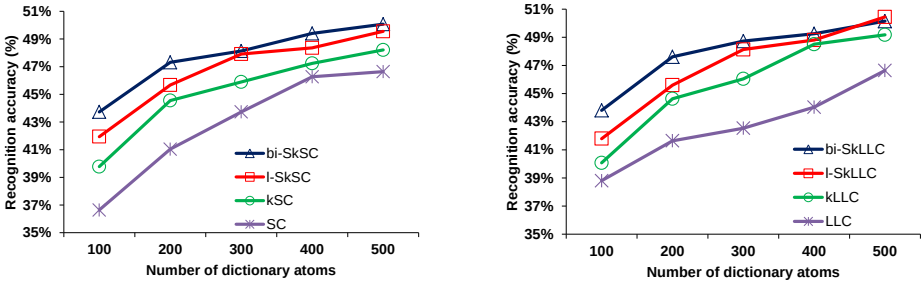

Fig. 3: mAP curves for dictionary learning with unsupervised and supervised kernel coding on MIT Indoor.

8000 visual words computed from dense RootSIFT [31] as initial image descriptor. The final descriptor was obtained by whitening and reducing the dimensionality of the initial descriptor to 4000. We followed the test protocol used in [33], where each category contains about 80 training images and 20 test images, and report the mean Average Precision (mAP).

In Fig. 3, we compare the performance of linear coding against kernel coding (supervised and unsupervised). Several observations can be made from Fig. 3. First, kernel coding, *i.e.*, kSC and kLLC, significantly outperform the linear methods, SC and LLC. For example, kLLC with a dictionary of size 500 reaches a mAP of 49.2% while LLC yields 46.6%. Second, supervised learning consistently improves the performance. The maximum mAP of 50.4% is achieved by l-SkLLC. Importantly, the results obtained here are competitive to several state-of-the-art techniques. For example, the mAP reported in [34] and [35] are of 49.4% and 50.15%, respectively. Given the simplicity of the image descriptor employed here, one could expect higher performance if more discriminative features (*e.g.*, Fisher vectors [36]) were used. This, however, goes beyond the scope of the paper, which rather aims to assess the performance of kernel coding against linear coding.

4.3 Pascal VOC2007

As a third experiment, we evaluated the proposed kernel coding methods on the PASCAL Visual Object Classes Challenge (VOC2007) [37]. This dataset contains 9963 images from 20 classes, including people, animals, vehicles and indoor objects, and is considered as a realistic and difficult dataset for object classification. We used the same descriptors as for the MIT Indoor dataset to represent the images. The dataset is split into training (2501 images), validation (2501 images) and test (4952 images) sets. We used the validation set to obtain the bandwidth of the Gaussian kernels used in our experiments.

In Fig. 4, we compare the performance of linear coding against kernel coding (supervised and unsupervised) for VOC2007. This figure reveals that

1. Kernel coding is superior to linear methods for all dictionary sizes. For example, kSC outperforms SC by almost 5% using 500 atoms.

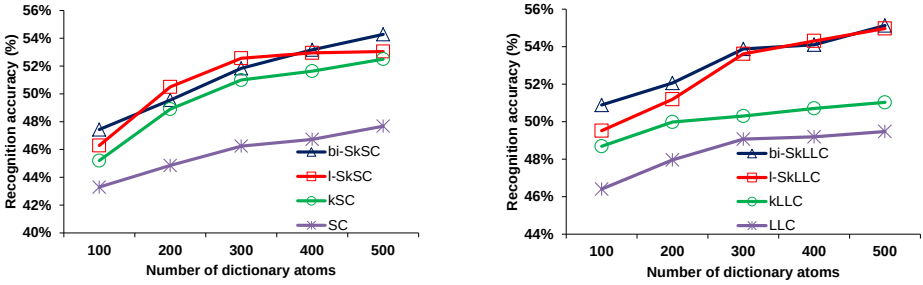

Fig. 4: Recognition accuracies of dictionary learning for unsupervised and supervised kernel coding VOC2007.

2. Supervised coding boosts the performance even further. For example, with the linear classifier from (23), the performance of kLLC can be improved from 50.7% to 54.3% with 400 atoms.
3. In the majority of cases, the bilinear supervised model outperforms the linear classifier. This, however, is not always the case. For example, with 300 atoms, the performance of l-SkSC is 52.6% while bi-SkSC yields 51.8%.

4.4 Extended YaleB

As a fourth experiment, we evaluated kernel coding against SRC [32] and the state-of-the-art LC-kSVD [19], which can be considered as a supervised extension of linear sparse coding and dictionary learning. To provide a fair comparison, we used the data provided by the authors of [19] for the extended YALE-B dataset [38]. The extended YaleB database contains 2,414 frontal face images of 38 people (*i.e.*, about 64 images for each person). We used half of the images per category as training data and the other half for testing purpose using the partition provided by [19].

We performed two tests on YALEB. In the first one, we evaluated the performance of kSC and kLLC by learning a dictionary of size 570 in an unsupervised manner. The results of this test are shown at the top of Table 2. We observe that both kSC and kLLC, without supervised learning, outperform the state-of-the-art LC-KSVD method. The maximum accuracy of 97.2% is achieved by kLLC which outperforms LC-KSVD by more than 2%. For the second test, we considered all training data as dictionary atoms and employed the classification method described in Eq.27. The results of this test are shown in the bottom part of Table 2. Again we see that kLLC achieves the highest accuracy of 98.4% and thus outperforms LC-KSVD.

4.5 Virus

Finally, to illustrate the ability of kernel coding methods at handling manifold data, we performed coding on a Riemannian manifold using the virus dataset [39]. The virus dataset contains 15 different virus classes. Each class comprises 100 images of size 41×41 that were automatically segmented [39]. We used the 10 splits provided with the dataset in a leave one out manner, *i.e.*, 10 experiments with 9 splits for training and

Method	Recognition Accuracy
SRC-(570 atoms)	80.5%
LC-KSVD-(570 atoms)	95.0%
kSC-(570 atoms)	96.9%
kLLC-(570 atoms)	**97.2**%
SRC-(all training samples)	97.2%
LC-KSVD-(all training samples)	96.7%
kSC-(all training samples)	98.2%
kLLC-(all training samples)	**98.4**%

Table 2: Recognition accuracies for the extended YaleB dataset [38].

1 split as query. In this experiment, we used Region Covariance Matrices (RCM) as image descriptors. To generate the RCM descriptor of an image, at each pixel (u, v) of the image, we computed the 25-dimensional feature vector

$$\boldsymbol{x}_{u,v} = \left[I_{u,v}, \left|\frac{\partial I}{\partial u}\right|, \left|\frac{\partial I}{\partial v}\right|, \left|\frac{\partial^2 I}{\partial u^2}\right|, \left|\frac{\partial^2 I}{\partial v^2}\right|, |G^{0,0}_{u,v}|, \cdots, |G^{4,5}_{u,v}| \right]^T,$$

where $I_{u,v}$ is the intensity value, $G^{o,s}_{u,v}$ is the response of a 2D Gabor wavelet [40] with orientation o and scale s, and $|\cdot|$ denotes the magnitude of a complex value. Here, we generated 20 Gabor filters at 4 orientations and 5 scales. RCM descriptors are symmetric positive definite matrices, and thus lie on a Riemannian manifold where Euclidean methods do not apply. Our framework, however, lets us employ the log-Euclidean RBF kernel recently proposed in [41]. With this kernel, kSC and kLLC achieved 79.1% and 80.0, respectively. In [39], the performance of various features, specifically designed for the task of virus classification, was studied. The state-of-the-art accuracy reported in [39] was 54.5%, which both of kernel coding methods clearly outperform. Aside from showing the power of kernel coding, this experiment also demonstrates that kernel coding can be employed in scenarios where linear coding is not applicable.

5 Conclusions and Future Work

In this paper, we have introduced a general formulation for coding in possibly infinite dimensional Reproducing Kernel Hilbert Spaces (RKHS). Among others, we have studied the kernelization of two popular coding schemes, namely sparse coding and locality-constrained linear coding, and have introduce several extensions such as supervised coding using linear or bilinear classifiers. Furthermore, we have proposed two ways to identify a discriminative RKHS for coding. Our experimental evaluation on several recognition tasks has demonstrated the benefits of the proposed kernel coding schemes over conventional linear coding solutions. In the future, we plan to investigate the notion of max-margin for supervised kernel coding. We are also interested in devising kernel extensions of other coding schemes such as Vector of Locally Aggregated Descriptors (VLAD) [42] and Fisher vectors [36].

References

1. Boureau, Y.L., Bach, F., LeCun, Y., Ponce, J.: Learning mid-level features for recognition. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2010) 2559–2566
2. Coates, A., Ng, A.Y.: The importance of encoding versus training with sparse coding and vector quantization. In: Proc. Int. Conference on Machine Learning (ICML). (2011) 921–928
3. Koenderink, J., Van Doorn, A.: The structure of locally orderless images. Int. Journal of Computer Vision (IJCV) **31**(2-3) (1999) 159–168
4. Grauman, K., Darrell, T.: The pyramid match kernel: Discriminative classification with sets of image features. In: Proc. Int. Conference on Computer Vision (ICCV). Volume 2., IEEE (2005) 1458–1465
5. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Volume 2. (2006) 2169–2178
6. Agarwal, A., Triggs, B.: Hyperfeatures–multilevel local coding for visual recognition. In: Proc. European Conference on Computer Vision (ECCV). Springer (2006) 30–43
7. van Gemert, J.C., Veenman, C.J., Smeulders, A.W., Geusebroek, J.M.: Visual word ambiguity. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(7) (2010) 1271–1283
8. Liu, L., Wang, L., Liu, X.: In defense of soft-assignment coding. In: Proc. Int. Conference on Computer Vision (ICCV), IEEE (2011) 2486–2493
9. Olshausen, B.A., et al.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature **381**(6583) (1996) 607–609
10. Yang, J., Yu, K., Gong, Y., Huang, T.: Linear spatial pyramid matching using sparse coding for image classification. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2009) 1794–1801
11. Yu, K., Zhang, T., Gong, Y.: Nonlinear learning using local coordinate coding. In: Proc. Advances in Neural Information Processing Systems (NIPS). Volume 9. (2009) 1
12. Wang, J., Yang, J., Yu, K., Lv, F., Huang, T., Gong, Y.: Locality-constrained linear coding for image classification. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2010) 3360–3367
13. Winn, J., Criminisi, A., Minka, T.: Object categorization by learned universal visual dictionary. In: Proc. Int. Conference on Computer Vision (ICCV). Volume 2. (2005) 1800–1807
14. Perronnin, F.: Universal and adapted vocabularies for generic visual categorization. IEEE Transactions on Pattern Analysis and Machine Intelligence **30**(7) (2008) 1243–1256
15. Lazebnik, S., Raginsky, M.: Supervised learning of quantizer codebooks by information loss minimization. IEEE Transactions on Pattern Analysis and Machine Intelligence **31**(7) (2009) 1294–1309
16. Mairal, J., Bach, F., Ponce, J., Sapiro, G., Zisserman, A., et al.: Supervised dictionary learning. In: Proc. Advances in Neural Information Processing Systems (NIPS). (2009) 1033–1040
17. Yang, J., Yu, K., Huang, T.: Supervised translation-invariant sparse coding. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2010) 3517–3524
18. Lian, X.C., Li, Z., Lu, B.L., Zhang, L.: Max-margin dictionary learning for multiclass image categorization. In: Proc. European Conference on Computer Vision (ECCV). Springer (2010) 157–170
19. Jiang, Z., Lin, Z., Davis, L.: Label consistent k-svd: Learning a discriminative dictionary for recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **35**(11) (2013) 2651–2664

20. Gao, S., Tsang, I.W.H., Chia, L.T.: Kernel sparse representation for image classification and face recognition. In: Proc. European Conference on Computer Vision (ECCV). Springer (2010) 1–14
21. Van Nguyen, H., Patel, V., Nasrabadi, N., Chellappa, R.: Design of non-linear kernel dictionaries for object recognition. Image Processing, IEEE Transactions on **22**(12) (Dec 2013) 5123–5135
22. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. Information processing & management **24**(5) (1988) 513–523
23. Elad, M.: Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing. Springer (2010)
24. Liu, J., Ji, S., Ye, J.: SLEP: Sparse Learning with Efficient Projections. Arizona State University. (2009)
25. Mairal, J., Bach, F., Ponce, J., Sapiro, G.: Online learning for matrix factorization and sparse coding. Journal of Machine Learning Research (JMLR) **11** (2010) 19–60
26. Boyd, S.P., Vandenberghe, L.: Convex optimization. Cambridge university press (2004)
27. Bucak, S., Jin, R., Jain, A.: Multiple kernel learning for visual object recognition: A review. IEEE Transactions on Pattern Analysis and Machine Intelligence (in press)
28. Schölkopf, B., Herbrich, R., Smola, A.J.: A generalized representer theorem. In: Computational learning theory, Springer (2001) 416–426
29. Sharan, L., Rosenholtz, R., Adelson, E.: Material perception: What can you see in a brief glance? Journal of Vision **9**(8) (2009) 784–784
30. Sharan, L., Liu, C., Rosenholtz, R., Adelson, E.H.: Recognizing materials using perceptually inspired features. Int. Journal of Computer Vision (IJCV) **103**(3) (2013) 348–371
31. Arandjelovic, R., Zisserman, A.: Three things everyone should know to improve object retrieval. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2012) 2911–2918
32. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. Pattern Analysis and Machine Intelligence, IEEE Transactions on **31**(2) (2009) 210–227
33. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2009) 413–420
34. Singh, S., Gupta, A., Efros, A.A.: Unsupervised discovery of mid-level discriminative patches. In: Proc. European Conference on Computer Vision (ECCV). Springer (2012) 73–86
35. Wang, X., Wang, B., Bai, X., Liu, W., Tu, Z.: Max-margin multiple-instance dictionary learning. In: Proc. Int. Conference on Machine Learning (ICML). (2013) 846–854
36. Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J.: Image classification with the fisher vector: Theory and practice. Int. Journal of Computer Vision (IJCV) **105**(3) (2013) 222–245
37. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. Int. Journal of Computer Vision (IJCV) **88**(2) (2010) 303–338
38. Georghiades, A.S., Belhumeur, P.N., Kriegman, D.: From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE Transactions on Pattern Analysis and Machine Intelligence **23**(6) (2001) 643–660
39. Kylberg, G., Sintorn, I.M.: Evaluation of noise robustness for local binary pattern descriptors in texture classification. EURASIP Journal on Image and Video Processing **2013**(1) (2013) 1–20
40. Lee, T.S.: Image representation using 2d Gabor wavelets. IEEE Transactions on Pattern Analysis and Machine Intelligence **18**(10) (1996) 959–971
41. Jayasumana, S., Hartley, R., Salzmann, M., Li, H., Harandi, M.: Kernel methods on the riemannian manifold of symmetric positive definite matrices. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2013) 73–80

42. Jégou, H., Perronnin, F., Douze, M., Sánchez, J., Pérez, P., Schmid, C.: Aggregating local image descriptors into compact codes. IEEE Transactions on Pattern Analysis and Machine Intelligence **34**(9) (2012) 1704–1716